

The Annals of Applied Statistics
 2012, Vol. 6, No. 4, 1795–1813
 DOI: 10.1214/12-AOAS569
 © Institute of Mathematical Statistics, 2012

LATENT DEMOGRAPHIC PROFILE ESTIMATION IN HARD-TO-REACH GROUPS

BY TYLER H. MCCORMICK¹ AND TIAN ZHENG²

University of Washington and Columbia University

The sampling frame in most social science surveys excludes members of certain groups, known as *hard-to-reach groups*. These groups, or subpopulations, may be difficult to access (the homeless, e.g.), camouflaged by stigma (individuals with HIV/AIDS), or both (commercial sex workers). Even basic demographic information about these groups is typically unknown, especially in many developing nations. We present statistical models which leverage social network structure to estimate demographic characteristics of these subpopulations using Aggregated relational data (ARD), or questions of the form “How many X’s do you know?” Unlike other network-based techniques for reaching these groups, ARD require no special sampling strategy and are easily incorporated into standard surveys. ARD also do not require respondents to reveal their own group membership. We propose a Bayesian hierarchical model for estimating the demographic characteristics of hard-to-reach groups, or *latent demographic profiles*, using ARD. We propose two estimation techniques. First, we propose a Markov-chain Monte Carlo algorithm for existing data or cases where the full posterior distribution is of interest. For cases when new data can be collected, we propose guidelines and, based on these guidelines, propose a simple estimate motivated by a missing data approach. Using data from McCarty et al. [*Human Organization* **60** (2001) 28–39], we estimate the age and gender profiles of six hard-to-reach groups, such as individuals who have HIV, women who were raped, and homeless persons. We also evaluate our simple estimates using simulation studies.

1. Introduction. Standard surveys often exclude members of the certain groups, known as *hard-to-reach groups*. One reason these individuals are excluded is difficulty accessing group members. Persons who are homeless are

Received February 2011; revised February 2012.

¹Supported in part by NIAID Grant R01 HD54511. This work was partially completed while Tyler McCormick was supported by a Google Ph.D. fellowship in Statistics.

²Supported by NSF Grant SES-1023176 and a Google faculty research award.

Key words and phrases. Aggregated relational data, hard-to-reach populations, hierarchical model, social network, survey design.

This is an electronic reprint of the original article published by the Institute of Mathematical Statistics in *The Annals of Applied Statistics*, 2012, Vol. 6, No. 4, 1795–1813. This reprint differs from the original in pagination and typographic detail.

very unlikely to be reached by a survey which uses random-digit dialing, for example. Other individuals can be accessed using standard survey techniques, but are excluded because of issues in reporting. Members of these groups are often reluctant to self-identify because of social pressure or stigma [Shelley et al. (1995)]. Individuals who are homosexual, for example, may not be comfortable revealing their sexual preferences to an unfamiliar survey enumerator. A third group of individuals is difficult to reach because of issues with both access and reporting (commercial sex workers, e.g.).

Even basic demographic information about these groups is typically unknown, especially in developing nations. We propose a Bayesian hierarchical model for estimating the demographic characteristics of hard-to-reach groups, or *latent demographic profiles*. Specifically, these profiles reveal features such as the number of males in a certain age range, say, 20–30 years old, who have HIV. Sociologically, this information yields insights into the characteristics of some of the most socially isolated members for the population. Along with its contribution to our understanding of contemporary social institutions, estimating demographic profiles for these groups also has public health benefits. The distribution of infected individuals influences the size of the public health response. UNAIDS—the joint United Nations program on HIV/AIDS, for example, currently sponsors several projects using a variety of techniques to estimate the sizes of populations most at-risk for HIV/AIDS [UNAIDS (2003)]. The proposed method would, along with estimating the size of the population, provide latent demographic profiles. This information would not only help calibrate the scale of the response but also tailor programs to the specific needs of population members.

One approach to estimating demographic information about hard-to-reach groups is to reach members of these groups through their social network. Some network-based approaches, such as Respondent-driven Sampling (RDS), recruit respondents directly from other respondents' networks [Heckathorn (1997, 2002)], making the sampling mechanism similar to a stochastic process on the social network [Goel and Salganik (2009)]. RDS affords researchers face-to-face contact with members of hard-to-reach groups, facilitating exhaustive interviews and even genetic or medical testing. The price for an entrée to these groups is high, however, as RDS uses a specially designed link-tracing framework for sampling. Estimates from RDS are also biased because of the network structure captured during selection, with much statistical work surrounding RDS being intended to reweight observations from RDS to have properties resembling a simple random sample.

Another approach is Aggregated relational data (ARD) or “How many X’s do you know” questions [Killworth et al. (1998a)]. In these questions, “X” defines a population of interest (e.g., How many people who are homeless do you know?). A specific definition of “know” defines the network the respondent references when answering the question. In contrast to RDS,

ARD do not require reaching members of the hard-to-reach groups directly. Instead, ARD access hard-to-reach groups indirectly through the social networks of respondents on standard surveys. ARD never affords direct access to members of hard-to-reach populations, making the level of detail achievable though RDS impossible with ARD. Unlike RDS, however, ARD require no special sampling techniques and are easily incorporated into standard surveys. ARD are, therefore, feasible for a broader range of researchers across the social sciences, public health, and epidemiology to implement with significantly lower cost than RDS.

In this paper, we propose a model for estimating latent demographic profiles using ARD. The ease of implementation of ARD means that the models proposed here will make the demographic characteristics of hard-to-reach groups available to the multitude of researchers collecting data using standard survey methodology. Specifically, we propose a Bayesian hierarchical model for estimating the demographic characteristics of hard-to-reach groups using ARD. When the full posterior is of interest, we propose a Markov-chain Monte Carlo algorithm.

Given the ease of collecting ARD, we speculate that many researchers may be interested in including ARD questions on future surveys. In this case, we show that estimates for some network features very close to those achieved using MCMC can be obtained using significantly simpler estimation techniques under certain survey design conditions. Along with giving survey guidelines, we propose a simpler estimation technique based on the EM algorithm and regression. Using data from McCarty et al. (2001), we estimate the age and gender profiles of six hard-to-reach groups, such as individuals who have HIV, women who were raped, and homeless persons.

In Section 2 we contextualize our proposed method by reviewing previous statistical methods for estimating network features using ARD. Then, we describe a method for estimating demographic profiles from hard-to-reach populations. Section 4 illustrates our method using data from McCarty et al. (2001). After demonstrating the utility of our model, Section 5 describes how, under certain survey design conditions, we can obtain similar estimates without the computational sophistication required by MCMC.

2. Previous research on ARD. ARD are commonly used to estimate the size of populations that are difficult to count directly. The scale-up method, an early method for ARD, uses ARD questions where the subpopulation size is known (people named Michael, e.g.) to estimate degree in a straightforward manner. Suppose that you know two persons named Nicole, and that at the time of the survey, there were 358,000 Nicos out of 280 million Americans. Thus, your two Nicos represent a fraction ($2/358,000$) of all the Nicos. Extrapolating to the entire country yields an estimate of $(2/358,000) \times (280 \text{ million}) = 1560$ people known by you. Then, the size of

unknown subpopulations is estimated by solving the given equation for the unknown subpopulation size with the estimated degree. Using this method, ARD has been used extensively to estimate the size of populations such as those with HIV/AIDS, injection drug users, or the homeless [e.g., Killworth et al. (1990, 1998b)].

The scale-up method is easy to implement but does not account for network structure. Consider, for example, asking a respondent how many people named “Rose” she/he knows. If knowing someone named Rose were entirely random, then each respondent would be equally likely to know each of the one-half million Rose’s on the hypothetical list; that is, each respondent on each Rose is a Bernoulli trial with a fixed success probability. Network structure makes these types of independence assumptions invalid. Since Rose is most common among older females and people are more likely to know individuals of similar age and the same gender, older female respondents are more likely to know a given Rose than older male respondents. Assuming independent responses induces bias in the individuals’ responses. Since estimates of hard-to-count populations are then constructed using responses to ARD, the resulting estimates are also biased [Killworth et al. (1998a), Bernard et al. (1991)].

Zheng, Salganik and Gelman (2006) and McCormick, Salganik and Zheng (2010) propose hierarchical models for ARD which partially address the manifestations of network structure present in ARD. McCormick, Salganik and Zheng (2010) develop a model specifically for estimating respondents’ degree (network size) and population degree distribution. Though this model accounts for the network structure described in the above example, McCormick, Salganik and Zheng (2010) do not address hard-to-reach groups. Zheng, Salganik and Gelman (2006) present a model which estimates the sizes of hard-to-reach groups [see Figure 5 in Zheng, Salganik and Gelman (2006)]. This paper presents a model which provides richer information about hard-to-reach groups by estimating both subpopulation sizes and the demographic breakdown of individuals within these groups.

3. Estimating latent profiles. In this section we describe a model for estimating latent demographic profiles for hard-to-reach groups. This method will provide information about the demographic makeup of groups which are often difficult to access using standard surveys, such as the proportion of young males who are infected with HIV. The observations, y_{ik} , represent the number of individuals in subpopulation k known by respondent i . In ARD, respondents are conceptualized as *egos*, or senders of ties in the network. We divide the egos into groups based on their demographic characteristics (males 20–40 years old, e.g.). The individuals who comprise the counts for ARD are the *alters*, or recipients of links in the network. The alters are also divided into groups, though the groups need not be the same for both

the ego and the alter groups. Under this setup members of hard-to-reach groups are one type of alter. Thus, determining the alter groups determines the demographic characteristics of the hard-to-reach groups which can be estimated. We model the number of people that respondent i is connected to in subpopulation k as

$$(3.1) \quad y_{ik} \sim \text{Neg-Binom}(\mu_{ike}, \omega_k)$$

$$\text{where } \mu_{ike} = d_i \sum_{a=1}^A m(e, a) h(a, k)$$

and ω_k represents the variation in the relative propensity of respondents within an ego group to form ties with individuals in a particular subpopulation k . The degree of person i is d_i and e is the ego group that person i belongs to. The $h(a, k)$ term is the relative size of subpopulation group k within alter group a (e.g., 4% of males between ages 21 and 40 are named Michael). The mixing coefficient, $m(e, a)$, for a respondent with degree $d_i = \sum_{a=1}^A d_{ia}$ between ego-group e and alter-group a is

$$m(e, a) = E\left(\frac{d_{ia}}{d_i} \mid i \text{ in ego group } e\right),$$

where d_{ia} is the number of person i 's acquaintances in alter group a . That is, $m(e, a)$ represents the expected fraction of the ties of someone in ego-group e that go to people in alter-group a . For any group e , $\sum_{a=1}^A m(e, a) = 1$.

The vector of mixing rates for an ego group, $(m(e, 1), \dots, m(e, A))^T$, enters the likelihood via an inner product with $h(a, k)$; therefore, its components are only identifiable if the A by K matrix of $h(a, k)$ terms, $\mathbf{H}_{A \times K}$, has rank A . This condition requires $K > A$ and that the columns of $\mathbf{H}_{A \times K}$ not be perfectly correlated. When all elements of $\mathbf{H}_{A \times K}$ are fixed, then (3.1) is the LRNM model from McCormick, Salganik and Zheng (2010). Specifically, McCormick, Salganik and Zheng (2010) propose asking ARD questions about populations where the elements of $h(a, k)$ are readily available, such as first names in the United States population. When $h(a, k)$ is known, it is simply N_{ak}/N_a or the number of individuals in alter group a who have characteristic k divided by the number of individuals in alter group a .

In hard-to-reach groups, $h(a, k)$ is rarely known. In many cases, even the number of individuals in a hard-to-reach group, N_k , is unknown. In the following section we propose a method for estimating $h(a, k)$ for hard-to-reach groups using information from groups when $h(a, k)$ is available. This method provides information beyond the size of the subpopulation group, also estimating the number of individuals in each of the a alter groups, N_{ak} .

In summary, the number of people that person i knows in subpopulation k , given that person i is in ego-group e , is based on person i 's degree (d_i), the proportion of people in alter-group a that belong to subpopulation k ,

$(h(a, k))$, and the mixing rate between people in group e and people in group a , $(m(e, a))$. Additionally, if we observe random mixing, then $m(e, a) = N_a/N$.

Similar to Zheng, Salganik and Gelman (2006), a negative binomial model is assumed in (3.1) for each y_{ik} with an overdispersion, ω_k , parameter measuring the residual relative propensity of respondents to form ties with individuals in group k , controlling for the variations that define the ego groups.

3.1. Latent demographic profiles. We propose a two-stage estimation procedure. We first use a multilevel model and Bayesian inference to estimate d_i , $m(e, a)$, and ω'_k using the latent nonrandom mixing model described in McCormick, Salganik and Zheng (2010) for the subpopulations where $h(a, k) = N_{ak}/N_a$ is known. Second, conditional on this information, we estimate the latent profiles for the remaining subpopulations.

For the estimation of the McCormick, Salganik and Zheng (2010) model components, we assume that $\log(d_i)$ follows a normal distribution with mean μ_d and standard deviation σ_d . Zheng, Salganik and Gelman (2006) postulate that this prior should be reasonable based on previous work, specifically McCarty et al. (2001), and found that the sampler described using this prior mixed well and satisfied posterior predictive checks. McCormick, Salganik and Zheng (2010) also conducted simulation experiments which demonstrated that the shape of the posterior was not an artifact of this prior assumption. We estimate a value of $m(e, a)$ for all E ego groups and all A alter groups. For ego group, e , and alter group, a , we assume that $m(e, a)$ has a normal prior distribution with mean $\mu_{m(e,a)}$ and standard deviation $\sigma_{m(e,a)}$. For ω'_k , we use independent uniform(0, 1) priors on the inverse scale, $p(1/\omega'_k) \propto 1$. Since ω'_k is constrained to $(1, \infty)$, the inverse falls on $(0, 1)$. The Jacobian for the transformation is ω'^{-2} . For the latent profiles, define $\mathbb{1}_{h(a,k)}$ as the indicator of the latent profiles. The matrix $h(a, k)$ is defined as N_{ak}/N_a when population information is available ($\mathbb{1}_{h(a,k)} = 0$) and entries to be estimated ($\mathbb{1}_{h(a,k)} = 1$) are given normal priors on the log scale with mean μ_h and standard deviation σ_h . That is, we model each $\log(h(a, k)) \sim N(\mu_h, \sigma_h^2)$ with a common mean and variance for all entries in the latent profile matrix. Since many of the profiles are close to zero, we found that the additional structure from a common prior across all entries improved convergence without being too rigid to capture fluctuations in latent intensity. Finally, we give noninformative uniform priors to the hyperparameters μ_d , $\mu_{m(e,a)}$, μ_h , σ_d and $\sigma_{m(e,a)}$, σ_h . The joint posterior density can then be expressed as

$$\begin{aligned} p(d, m(e, a), \omega', \mu_d, \mu_{m(e,a)}, \sigma_d, \sigma_{m(e,a)} | y) \\ \propto \prod_{k=1}^K \prod_{i=1}^N \binom{y_{ik} + \xi_{ik} - 1}{\xi_{ik} - 1} \left(\frac{1}{\omega'_k}\right)^{\xi_{ik}} \left(\frac{\omega'_k - 1}{\omega'_k}\right)^{y_{ik}} \end{aligned}$$

$$\begin{aligned}
& \times \left(\frac{1}{\omega_k} \right)^2 \prod_{i=1}^N N(\log(d_i) | \mu_d, \sigma_d) \\
& \times \prod_{e=1}^E N(m(e, a) | \mu_{m(e, a)}, \sigma_{m(e, a)}) \\
& \times \mathbb{1}_{h(a, k)} \prod_{k=1}^K \prod_{a=1}^A N(h(a, k) | \mu_{h(a, k)}, \sigma_{h(a, k)}),
\end{aligned}$$

where $\xi_{ik} = d_i f(\sum_{a=1}^A m(e, a) h(a, k)) / (\omega'_k - 1)$.

Adapting Zheng, Salganik and Gelman (2006) and McCormick, Salganik and Zheng (2010), we use a Gibbs–Metropolis algorithm in each iteration v :

- (1) For each i , update d_i using a Metropolis step with jumping distribution $\log(d_i^*) \sim N(d_i^{(v-1)}, (\text{jumping scale of } d_i)^2)$.
- (2) For each e , update the vector $m(e, \cdot)$ using a Metropolis step. Define the proposed value using a random direction and jumping rate. Each of the A elements of $m(e, \cdot)$ has a marginal jumping distribution $m(e, a)^* \sim N(m(e, a)^{(v-1)}, (\text{jumping scale of } m(e, \cdot))^2)$. Then, rescale so that the row sum is one.
- (3) Update $\mu_d \sim N(\hat{\mu}_d, \sigma_d^2/n)$, where $\hat{\mu}_d = \frac{1}{n} \sum_{i=1}^n d_i$.
- (4) Update $\sigma_d^2 \sim \text{Inv-}\chi^2(n-1, \hat{\sigma}_d^2)$, where $\hat{\sigma}_d^2 = \frac{1}{n} \times \sum_{i=1}^n (d_i - \mu_d)^2$.
- (5) Update $\mu_{m(e, a)} \sim N(\hat{\mu}_{m(e, a)}, \sigma_{m(e, a)}^2/A)$, for each e where $\hat{\mu}_{m(e, a)} = \frac{1}{A} \sum_{a=1}^A m(e, a)$.
- (6) Update $\sigma_{m(e, a)}^2 \sim \text{Inv-}\chi^2(A-1, \hat{\sigma}_{m(e, a)}^2)$, for each e where $\hat{\sigma}_{m(e, a)}^2 = \frac{1}{A} \times \sum_{a=1}^A (m(e, a) - \mu_{m(e, a)})^2$.
- (7) For each k with a known profile, update ω'_k using a Metropolis step with jumping distribution $\omega'^*_k \sim N(\omega'^{(v-1)}_k, (\text{jumping scale of } \omega'_k)^2)$.

We now proceed to estimate the H latent profiles:

- (8) For each element of $h(a, k)$ where $\mathbb{1}_{h(a, k)} = 1$, update $h(a, k)$ using a Metropolis step with jumping distribution $h(a, k)^* \sim N(h(a, k)^{(v-1)}, (\text{jumping scale of } h(a, k))^2)$.
- (9) Update $\mu_h \sim N(\hat{\mu}_h, \sigma_h^2/(A \times H))$ for each k where

$$\hat{\mu}_h = \frac{1}{(A \times H)} \sum_{k=1}^K \sum_{a=1}^A \mathbb{1}_{h(a, k)} h(a, k).$$

- (10) Update $\sigma_h^2 \sim \text{Inv-}\chi^2((A \times H) - 1, \hat{\sigma}_h^2)$ where

$$\hat{\sigma}_h^2 = \frac{1}{A \times H} \sum_{k=1}^K \sum_{a=1}^A \mathbb{1}_{h(a, k)} (h(a, k) - \mu_h)^2.$$

- (11) For each k where $h(a, k)$ is estimated, update ω'_k using a Metropolis step with jumping distribution $\omega'^*_k \sim N(\omega'^{(v-1)}_k, (\text{jumping scale of } \omega'_k)^2)$.

Having $h(a, k)$ for some subpopulations is critical to estimating latent structure through latent profiles. Often, $h(a, k)$ can be obtained from publicly available sources (Census Bureau, Social Security Administration, etc.) for subpopulations such as first names. The number of populations with known $h(a, k)$ impacts the precision of the estimates for subpopulations with unknown profiles. Adding another known subpopulation increases the hypothetical sample size of each question, in essence asking each respondent if they know more alters. McCormick, Salganik and Zheng (2010) show that the total size of the subpopulations asked is related to the variance of estimated degree. Since known subpopulations are used to estimate degree, adding another subpopulation impacts variability in degree estimation in the first stage of our procedure, which propagates to estimates of $h(a, k)$. The alter groups where information is available for known $h(a, k)$ also limit the type of latent structure that can be estimated. McCormick, Salganik and Zheng (2010) create alter groups based on age and gender but note that separating alters based on other factors (such as race) would provide valuable information. The Census Bureau collects the information required to conduct such an analysis; however, McCormick, Salganik and Zheng (2010) report that their efforts to obtain the data were ultimately unsuccessful.

The choice of populations with known $h(a, k)$ is also important in ensuring that the mixing matrix is estimated appropriately. First, the subpopulations with known sizes need to be sufficiently heterogeneous with respect to their interactions with the ego groups to adequately estimate the mixing matrix. If, for example, our mixing matrix consists of only gender and we chose to use first names for subpopulations with known $h(a, k)$, then we should use a set of both male and female names. If we only asked male names, then we could estimate the propensity for males/females to interact with males but could not estimate the propensity of either gender to interact with females. Second, we make an assumption about the representativeness of respondents' networks rather than of the respondents themselves. For our method it would not be an issue, for example, if we recruit a smaller fraction of men into the survey than the proportion of men in the population. Instead, we would encounter bias if the networks of the men we selected were not representative of male networks in the population. This could happen, for example, if we recruit only men who know a disproportionately large number of women. This issue could also be exacerbated by differential nonresponse. Consider, for example, the case where individuals who know members of the hard-to-reach groups are less likely to answer questions than the general population. We continue this discussion in Section 4 where we postulate that errors in the estimates obtained in our data could be from bias in the estimates of the mixing matrix.

Finally, certain types of bias which are consistently associated with ARD should also be considered when selecting the subpopulations with known $h(a, k)$. We assume, for example, that the responses are free from transmission error, when a respondent knows a member of a subpopulation but is unaware of the alter's membership. McCormick, Salganik and Zheng (2010) suggest using first names since they represent the minimum conceivable possibility of transmission error. We also assume that respondents accurately recall the number of individuals they know in a given subpopulation. In reality, underestimation is common in large groups [see McCormick and Zheng (2007) for a detailed discussion].

4. Results for hard-to-count populations. We use data from a telephone survey by McCarty et al. (2001) with 1375 respondents and twelve names with known demographic profiles. These data have been analyzed in several previous studies and are typical ARD which are becoming increasingly common. The age and gender profiles of the names are available from the Social Security Administration. On this survey, “know” is defined “that you know them and they know you by sight or by name, that you could contact them, that they live within the United States, and that there has been some contact (either in person, by telephone, or mail) in the past 2 years.” We then estimate latent profiles for seven subpopulations. Six are groups often considered hard-to-count while the seventh uses ARD to learn about population social structure. Figure 1 displays the latent profiles for six populations which are often described as hard-to-count. For both individuals with HIV and those with AIDS we estimate the highest concentration to be among youth and young adult respondents. We estimate a higher concentration of young adult males than females for both HIV and AIDS with the concentration decreasing with age.

Subpopulations such as victims of homicide or persons who have committed suicide portray a key advantage of using ARD for measuring these populations. Our model estimates characteristics of these populations without requiring members of these populations to be reached directly through our survey. We compared our estimates of the number of individuals murdered in the past year with the 1999 *Uniform Crime Reports (UCR)* [Federal Bureau of Investigation (1999)] and figures from the Centers for Disease Control National Center for Injury Prevention and Control (CDC) [Centers for Disease Control (2011)]. A technical distinction between the two sources for external validation is that the CDC figures measure homicides (killing of another person) while UCR tally murders (unlawful killing of another person). So-called justifiable homicides (police officers using deadly force, e.g.) are therefore not counted in the UCR figures. This distinction accounts for part of the discrepancy between the two data sources (the FBI only keeps records on firearms-related justifiable homicides), though the exact amount

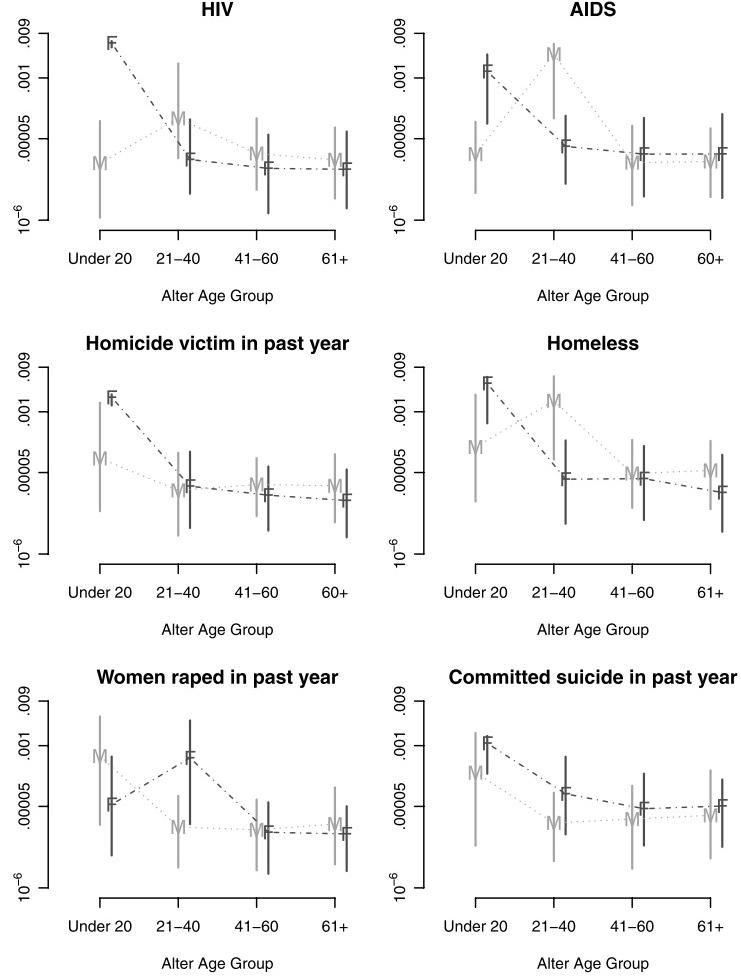


FIG. 1. Estimates of latent profiles for six hard-to-reach populations. The lighter text represents males and the darker text females. Letters correspond to posterior medians, while lines represent the width of the middle half of the posterior distribution. The estimated profiles are consistent with contemporary understanding of the profiles of these groups.

could not be determined from available data. Also, the McCarty et al. (2001) survey took place partially in January of 1999 and partially in June, meaning that this report does not capture precisely the period respondents were asked to recall. Since homicide statistics do not typically change drastically on a national scale over the course of a year, we expect, nonetheless, that these figures are reasonable for comparison. In all six age-gender categories, the UCR and CDC estimates are within the middle 50% of the posterior distribution of our estimates [computed by multiplying $h(a, k)$ by the number

of individuals in the given age-gender group]. For males 20–40, for example, the UCR counts approximately 5300 murders while that CDC counts just under 7300 homicides. Our method estimates the first quartile of the posterior distribution as roughly 300 murders and the third quartile as around 7300. Similarly, for females between 40 and 60 the middle half of our posterior lies between around 100 and 2300 while the UCR records around 700 and the CDC counts about 1900. Overall our estimates underrepresent the disparity in the proportion of male and female homicide victims, which we believe is due to the individuals who are most likely associated with murdered individuals being underrepresented in the survey frame. McCormick et al. (2009) found a similar issue in an internet survey.

Our estimates for women who were raped in the past year reveal a common issue with ARD questions. Though the questions asks respondents to recall only women who were raped, we hypothesize respondents will include men who are connected to a woman who was raped, even if the woman does not meet the definition of a tie. Respondents may also be likely to over-recall such traumatic events. Similarly, our estimates for female suicides are consistently higher than for males (though the difference is well within the uncertainty of measurements). Males are actually nearly four times as likely to commit suicide as females [Centers for Disease Control and Prevention (2011)]. This discrepancy might be because of the isolation of many suicide victims before their deaths, making them difficult to reach with ARD. Our estimates are especially consistent with the case of males being more isolated than females before committing suicide.

Recent work has used ARD for estimating population-level social phenomenon outside the context of hard-to-reach groups [DiPrete et al. (2011)]. To demonstrate the applicability of latent profile estimation in this context, Figure 2 shows the latent profile of individuals who opened a small business in the past year. The trend across ages in the profiles for males and females is similar, with most new business openers being younger adults [Office of Advocacy, U.S. Small Business Administration (1997)]. The fraction of males opening a business is consistently higher, however. This discrepancy is especially pronounced among young adults, the group with highest overall propensity.

Overall, our estimates of latent profiles are similar to estimates from other sources for the U.S. population. The similarity between previous knowledge about the profiles of these populations and our estimates indicates that ARD contain a significant amount of information about the latent structure of these subpopulations. The estimates presented in this section were obtained using the MCMC algorithm described in Section 3. In the following section we present an alternative regression-based estimation strategy which is significantly less time-consuming to implement and provides comparable performance when certain conditions are satisfied.

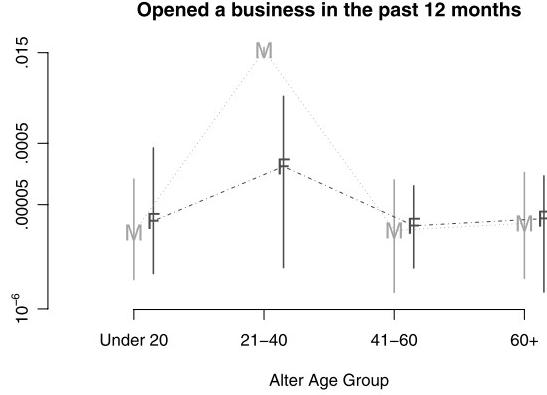


FIG. 2. *Estimates of latent profiles for individuals starting their own business. Letters correspond to posterior medians, while lines represent the width of the middle half of the posterior distribution. The lighter text represents males and the darker text females. Overall estimates are higher for males than for females, with the largest discrepancy for young adults.*

5. Simple calculations and design recommendations. Given data from an existing survey, we have shown that our method will recover features of unobserved subpopulation profiles. We propose an alternative strategy to recover this information under certain conditions without using MCMC. Our simple method combines estimation and survey-design strategy, making it well-suited for researchers who intend to collect ARD. McCormick, Salganik and Zheng (2010) proposed the *scaled-down* condition for selecting subpopulations to reduce bias in simple estimates of respondent degree. To estimate latent profiles, we need accurate degree and mixing matrix estimates. To accurately estimate the mixing matrix, we introduce a missing data perspective for ARD and propose an estimator based on the EM algorithm.

In Section 5.1 we review degree estimation and the scaled-down condition. Next, Section 5.2 describes a simple ratio estimator for the mixing matrix motivated by the EM algorithm. We then describe a regression based estimator for latent profiles in Section 5.3 and demonstrate its effectiveness through simulation studies in Section 5.4.

5.1. Estimating degree. In this section we review work on estimating respondent degree by McCormick, Salganik and Zheng (2010). We use these estimates in subsequent sections to estimate mixing rates and latent profiles.

McCormick, Salganik and Zheng (2010) develop a degree estimator based on the scale-up method of Killworth et al. (1998b). This approach uses respondents' answers to ARD questions and recalibrates based on the proportion of the total population comprised of the populations used on the survey.

For example, if a respondent reports knowing 3 women who gave birth, this represents about 1-millionth of all women who gave birth within the last year. This information then could be used to estimate that the respondent knows about 1-millionth of all Americans, $(3/3.6 \text{ million}) \cdot (300 \text{ million}) \approx 250$ people.

The precision of this estimate can be increased by averaging responses of many groups, yielding the scale-up estimator [Killworth et al. (1998b)],

$$\hat{d}_i = \frac{\sum_{k=1}^K y_{ik}}{\sum_{k=1}^K N_k} \cdot N,$$

where y_{ik} is the number of people that person i knows in subpopulation k , N_k is the size of subpopulation k , and N is the size of the population.

The scale-up estimator is easy to compute, yet can induce substantial bias if subpopulations aren't selected correctly. The scale-up estimator assumes *random mixing* across the K populations. That is, that the propensity for an individual to know members of a subpopulation depends only on the size of the subpopulation. In practice, this is rarely the case, as individuals tend to know more alters who are demographically similar to themselves.

McCormick, Salganik and Zheng (2010) derived a *scaled-down* condition for selecting names so that the collection of individuals with first names that are used to collect ARD constitute a balanced and representative sample of the population. In other words, the combined demographic profiles of the used first names match those of the general population. Specifically,

$$\frac{\sum_{k=1}^K N_{ak}}{N_a} = \frac{\sum_{k=1}^K N_k}{N} \quad \forall a.$$

Using the scaled-down condition, McCormick, Salganik and Zheng (2010) demonstrate that the scale-up estimator produces reduced-bias estimates of degree. In deriving the subsequent latent profile estimates, we assume we have selected subpopulations which satisfy the scaled-down condition.

5.2. A simple ratio estimator of individual mixing rates. If for a given respondent, i , we could take all the members of the social network with which i has a link and place them in a room, we would compute the mixing rate between the ego and a given alter group, $a = (1, \dots, A)$, by dividing the room in A mutually exclusive sections and asking alters to stand in their respective group. The estimated mixing rate would then be the number of people standing in a given group divided by the number of people in the room.

We could also perform a similar calculation by placing a simple random sample of size n from a population of size N in a room. Then, after dividing the alters into mutually exclusive groups, we could count y_{ia} or the number of alters respondent i knows in the sample who are in each of the a alter

groups. Since we have a simple random sample, we can extrapolate back to the population and estimate the degree of the respondent, \hat{d}_i , and within alter group degree, \hat{d}_{ia} , as

$$\hat{d}_i = \sum_{a=1}^A y_{ia}/(n/N) \quad \text{and} \quad \hat{d}_{ia} = y_{ia}/(n_a/N_a).$$

Given these two quantities, we can estimate the mixing rate between the respondent and an alter group by taking the ratio of alters known in the sample who are in alter group a over the total number known in the sample. This computation is valid because we assumed a simple random sample, thus that (in expectation) the demographic distribution of alters in our sample matches that of the population.

In ARD, the distribution of the hypothetical alters we sample depends on the subpopulations we select. If we only ask respondents subpopulations which consist of young males, for example, then our hypothetical room from the previous example would contain only the respondent's young, male alters. Estimating the rate of mixing between the respondent and older females would not be possible in this situation. Viewed in this light, ARD is a form of cluster sampling where the subpopulations are the clusters and respondents report the presence/absence of a tie between all alters in the cluster. Since the clusters are no longer representative of the population, our estimates need to be adjusted for the demographic profiles of the clusters [Lohr (1999)]. Specifically, if we observe y_{ika} for subpopulations $k = (1, \dots, K)$ and alter groups $a = (1, \dots, A)$, then our estimates of \hat{d}_i and \hat{d}_{ia} become

$$\hat{d}_i = \sum_{k=1}^K y_{ik} / \left(\sum_{k=1}^K N_k / N \right) \quad \text{and} \quad \hat{d}_{ia} = \sum_{k=1}^K y_{ika} / \left(\sum_{k=1}^K N_{ak} / N_a \right),$$

where N_k is the size of subpopulation k and N_{ak} is the number of members of subpopulation k in alter group a . To estimate the mixing rate, we could again divide the estimated number known in alter group a by the total estimated number known. Under the *scaled-down* condition the denominators in the above expressions cancel and the mixing estimate is the number known in the subpopulations that are in alter group a over the total number known in all K subpopulations.

In the examples above, we have assumed the alters are observed so that y_{ika} can be computed easily. This is not the case in ARD, however, since we observe only the aggregate number of ties and not the specific demographic makeup of the recipients. Thus, ARD are a cluster sample where the specific ties between the respondent and members of the alter group are *missing*.

If we ignore the residual variation in propensity to form ties with group k individuals due to noise [see (3.1) in Section 3], we may assume that the num-

ber of members of subpopulation k in alter group a the respondent knows, y_{ika} , follows a Poisson distribution. Under this assumption, we can estimate m_{ia} by imputing y_{ika} as part of an EM algorithm [Dempster, Laird and Rubin (1977)]. Specifically, for each individual define $\mathbf{y}_{ik}^{(\text{com})} = (y_{ika}, \dots, y_{i1A})^T$ as the complete data vector for each alter group. The complete data log-likelihood for individual i 's vector of mixing rates, $\mathbf{m}_i = (m_{i1}, \dots, m_{iA})^T$, is $\ell(\mathbf{m}_i; \mathbf{y}_{i1}^{(\text{com})}, \dots, \mathbf{y}_{iK}^{(\text{com})})$, which has the form

$$(5.1) \quad \begin{aligned} & \ell(\mathbf{m}_i; \mathbf{y}_{i1}^{(\text{com})}, \dots, \mathbf{y}_{iK}^{(\text{com})}) \\ &= \sum_{k=1}^K \sum_{a=1}^A \log \left(\text{Poisson} \left(y_{ika}; \lambda_{ika} = d_i m_{ia} \frac{N_{ak}}{N_a} \right) \right). \end{aligned}$$

Using (5.1), we derive the following two updating steps for the EM:

$$\begin{aligned} y_{ik}^{(t)} &= y_{ik} \left(\frac{m_{ia}^{(t-1)} (N_{ak}/N_a)}{\sum_{a=1}^A m_{ia}^{(t-1)} (N_{ak}/N_a)} \right), \\ m_{ia}^{(t)} &= \frac{\sum_{k=1}^K y_{ik}^{(t-1)}}{\sum_{k=1}^K y_{ik}}. \end{aligned}$$

If one sets $m_{ia}^{(0)} = N_a/N$, which corresponds to random mixing in the population, and runs one EM update, this would result in the following *simple ratio estimator* of the mixing rate for individual i :

$$(5.2) \quad \hat{m}_{ia} = \frac{\sum_{k=1}^K y_{ik} (N_{ak}/N_k)}{\sum_{k=1}^K y_{ik}}.$$

In our simulation studies, this simple estimator produces estimates very close to the converged EM estimates. Additionally, it is easy to show that the simple ratio estimate, \hat{m}_{ia} , is unbiased if $N_{ak}/N_a \neq 0$ for only one alter group a and that for any a there exists a subpopulation, k , such that $N_{ak} = N_a$. We refer to this condition as *complete separability*. Therefore, (5.2) constitutes a simple estimate for individual mixing rate and can be used to estimate average mixing behaviors of any ego group.

5.3. Regression-based estimates for latent profiles. The estimates for respondent degree and mixing estimates rely on latent profile information from some populations. Using these estimates, we now develop a regression-based estimator for unobserved latent profiles. For each respondent and each unknown subpopulation we now have

$$(5.3) \quad y_{ik} = \sum_{a=1}^A \hat{d}_i \hat{m}_{ia} h(a, k).$$

If we denote \mathbf{X}_k as the $n \times A$ matrix with elements $\hat{d}_i \hat{m}_{ia}$ and the vector $h(\cdot, k) = \vec{\beta}_k$, then (5.3) can be regarded as a linear regression equation, $\vec{y}_k = \mathbf{X}_k \vec{\beta}_k$, with the constraint that coefficients, $\vec{\beta}_k$, are restricted to be nonnegative. Lawson and Hanson (1974) propose an algorithm for computing these coefficients. Since the \hat{m}_{ik} sum to one across alter groups, the columns of \mathbf{X}_k are collinear. This could produce instability in solving the quadratic programming problem associated with finding our estimated latent profiles. In practice, we have found our estimates perform well despite this feature.

5.4. Simulation experiments. We present simulation experiments to evaluate our regression-based estimates under four strategies for selecting observed profiles. First, we created profiles which are completely separable. Second, we constructed profiles for the names satisfying the scaled-down condition presented in McCormick, Salganik and Zheng (2010) using data from the Social Security Administration. These names provide insights into the potential accuracy of our method using actual profiles. As a third case, we include the names from McCormick, Salganik and Zheng (2010) which violate the scaled-down condition and are almost exclusively popular among older respondents. For the fourth set of names, recall from Section 3 that the mixing matrix estimates are identifiable only if the matrix of known profiles, $\mathbf{H}_{A \times K}$, has rank A . To demonstrate a violation of this condition, we selected a set of names with uniform popularity across the demographic groups, or nearly perfect collinearity. There is some correlation in the scaled-down names since several names have similar profiles. The degree of correlation is substantially less than in the flat profiles, however.

In each simulation, we generated 500 respondents using the Latent Non-random Mixing Model in (3.1) [see also McCormick, Salganik and Zheng (2010)] with each of the four profile strategies. Mixing matrix estimates were calculated using the simple estimate derived from the first step of the EM algorithm in Section 5.2. We compare our mixing matrix estimates to the estimated mixing matrix from McCormick, Salganik and Zheng (2010), which we use to generate the simulated data. We evaluate the latent profiles using six names with profiles known from the Social Security Administration. We repeated the entire process 1000 times. Figure 3 presents boxplots of the squared error in mixing matrix and latent profile estimates. In both cases, the ideal, completely separable, profiles have the lowest error. The scaled-down names also perform well, indicating that reasonable estimates are possible even when complete separability is not. The flat profiles perform only slightly worse than the scaled-down names for estimating mixing but significantly worse when estimating latent profiles. The names which violate the scaled-down condition produce poor estimates of both quantities.

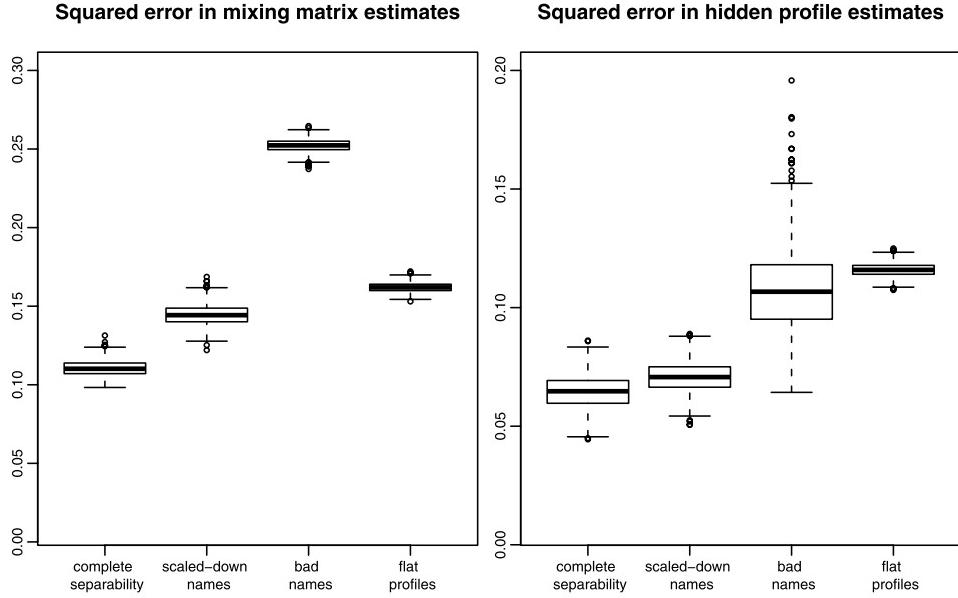


FIG. 3. Total mean squared error across all elements of the mixing matrix and latent profile matrix. The vertical axis is the sum of the errors across all eight alter groups. We generated 500 respondents using the four profile structures, then evaluated our ability to recover the mixing matrix estimated in McCormick, Salganik and Zheng (2010) and the known profiles of six additional names. We repeated the simulation 1000 times. In both cases the ideal profile has the lowest error, followed by the scaled-down names suggested by McCormick, Salganik and Zheng (2010).

6. Conclusion. We present a method for estimating latent profiles in hard-to-reach groups using standard surveys. Our method has two stages. First, we use known profiles for some populations to estimate respondent degree and the rate of mixing between survey respondents and groups in the population. Next, conditional on these estimates, we infer latent structure in populations where profiles are unknown. For existing data, we present a Bayesian hierarchical model and MCMC algorithm. We also propose viewing ARD in the context of missing data and provide a simple ratio estimate of mixing rates based on the EM algorithm. We then describe a regression-based estimate for latent profiles.

Despite its utility, there are several known issues with ARD. Using ARD in hard-to-reach populations presents special challenges which intersect with these known issues. Many events in this context are especially traumatic, leaving a more persistent signal in the respondent's memory than a typical tie. This phenomenon causes respondents to over-count their ties with a specific subpopulation. In Section 4 we contend that our overestimation of the proportion of men who are women who were raped in the past year is due to respondents overestimating by counting males who are associated

with females who have been raped, for example. This issue is in some sense the opposite of that faced by early ARD surveys for degree estimation when the concern was respondents under-recalling acquaintances from large populations [Killworth et al. (2003)]. Hard-to-reach groups are also often more open to interpretation than standard subpopulations. McCarty et al. (2001) give the example of people opening their own business and the homeless, for example. While there is some ambiguity in whether or not an individual has opened a new business, there is likely much greater variability between respondents in their classification of an individual as homeless. Hard-to-reach groups are also often associated with social stigma. This stigma increases the likelihood that a respondent will know a member of a subpopulation but not be aware that the alter belongs to the subpopulation, known as transmission errors. Recent work by Salganik et al. (2011) offers new insights into the magnitude of transmission errors in the context of HIV/AIDS, though the nature of the error likely depends heavily on the specific group of interests (respondents' decisions to reveal HIV status are likely quite different than their decision to discuss diabetes, e.g.).

This method also makes an assumption that the networks of the respondents are representative of networks of similar individuals in the population. In Section 4, in our discussion of the ratio of males to females who commit suicide, another possible explanation is that our survey does not include enough individuals who are likely to know people who commit suicide. This bias could be present in the networks of respondents even if the sample is, within respondents, observably representative. This point demonstrates the potential for future work in modeling bias that comes not from the respondents selected, but from the features of the networks of these respondents. This type of sampling bias is related to previous work by Lavallée (2007) and could prove a promising area for future work.

Our method demonstrates that ARD capture aspects of latent social structure through indirect observations of the social network. To do this, however, we require known profiles for some subpopulations. This requirement limits the estimable latent profiles to features which are known for some subpopulation. In our examples we use first names and estimate age and gender profiles. We may, for example, be interested in the race/ethnic profiles of the hard-to-reach populations. We are unable to estimate this from our current data because of the issues with obtaining demographic profiles for first names mentioned in Section 4. An alternative approach, and direction for potential future work, would be estimating a geometric, multidimensional latent social space based on features of the actors and the social network [Hoff, Raftery and Handcock (2002), Hoff (2005)]. Such a technique would provide a sense of the broad topography of the network (similar to Bayesian multi-dimensional scaling) and elucidate similarities between network structure in hard-to-reach groups.

An additional direction for future work involves combining information from ARD with other forms of data collection to better understand hard-to-reach groups. As mentioned in Section 1, RDS provides detailed information about a biased sample of members of the hard-to-reach group. This detailed information is in contrast to the indirect, general information obtained through ARD. The missing-data framework presented in Section 5 provides a first-step toward a general framework for combining information across various network-based data collection strategies.

Acknowledgments. The authors gratefully acknowledge the support of the National Science Foundation, the Columbia Applied Statistics Center, and the Columbia Population Research Center, along with the helpful comments of the Editor, Associate Editor, and an anonymous Referee.

REFERENCES

- BERNARD, H. R., JOHNSEN, E. C., KILLWORTH, P. D. and ROBINSON, S. (1991). Estimating the size of an average personal network and of an event subpopulation: Some empirical results. *Social Science Research* **20** 109–121.
- Centers for Disease Control (2011). WISQARS leading causes of death reports.
- Centers for Disease Control and Prevention (2011). Centers for Disease Control and Prevention, National Center for Injury Prevention and Control. Web-based Injury Statistics Query and Reporting System (WISQARS).
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **39** 1–38. [MR0501537](#)
- DiPRETE, T. A., GELMAN, A., MCCORMICK, T. H., TEITLER, J. and ZHENG, T. (2011). Segregation in social networks based on acquaintanceship and trust. *American Journal of Sociology* **116** 1234–1283.
- GOEL, S. and SALGANIK, M. J. (2009). Respondent-driven sampling as Markov chain Monte Carlo. *Stat. Med.* **28** 2202–2229. [MR2751515](#)
- HECKATHORN, D. D. (1997). Respondent-driven sampling: A new approach to the study of hidden populations. *Social Problems* **44** 174–199.
- HECKATHORN, D. D. (2002). Respondent-driven sampling II: Deriving valid population estimates from chain-referral samples of hidden populations. *Social Problems* **49** 11–34.
- HOFF, P. D. (2005). Bilinear mixed-effects models for dyadic data. *J. Amer. Statist. Assoc.* **100** 286–295. [MR2156838](#)
- HOFF, P. D., RAFTERY, A. E. and HANDCOCK, M. S. (2002). Latent space approaches to social network analysis. *J. Amer. Statist. Assoc.* **97** 1090–1098. [MR1951262](#)
- KILLWORTH, P. D., JOHNSEN, E. C., BERNARD, H. R., SHELLY, G. A. and McCARTY, C. (1990). Estimating the size of personal networks. *Social Networks* **12** 289–312.
- KILLWORTH, P. D., McCARTY, C., BERNARD, H. R., SHELLY, G. A. and JOHNSEN, E. C. (1998a). Estimation of seroprevalence, rape, and homelessness in the U.S. using a social network approach. *Evaluation Review* **22** 289–308.
- KILLWORTH, P. D., JOHNSEN, E. C., McCARTY, C., SHELLY, G. A. and BERNARD, H. R. (1998b). A social network approach to estimating seroprevalence in the United States. *Social Networks* **20** 23–50.

- KILLWORTH, P. D., McCARTY, C., BERNARD, H. R., JOHNSEN, E. C., DOMINI, J. and SHELLY, G. A. (2003). Two interpretations of reports of knowledge of subpopulation sizes. *Social Networks* **25** 141–160.
- LAVALLÉE, P. (2007). *Indirect Sampling*. Springer, New York. [MR2344875](#)
- LAWSON, C. L. and HANSON, R. J. (1974). *Solving Least Squares Problems*. Prentice Hall International, Englewood Cliffs, NJ. [MR0366019](#)
- LOHR, S. L. (1999). *Sampling: Design and Analysis*. Duxbury Press, Belmont, CA.
- MCCARTY, C., KILLWORTH, P. D., BERNARD, H. R., JOHNSEN, E. and SHELLY, G. A. (2001). Comparing two methods for estimating network size. *Human Organization* **60** 28–39.
- MCCORMICK, T. H., SALGANIK, M. J. and ZHENG, T. (2010). How many people do you know? Efficiently estimating personal network size. *J. Amer. Statist. Assoc.* **105** 59–70. [MR2757192](#)
- MCCORMICK, T. H. and ZHENG, T. (2007). Adjusting for recall bias in “How Many X’s Do You Know?” surveys. In *Proceedings of the Joint Statistical Meetings*. Salt Lake City, UT.
- MCCORMICK, T. H., MOUSSA, A., RUF, J., DiPRETE, T. A., GELMAN, A., TEITLER, J. and ZHENG, T. (2009). Comparing two methods for predicting opinions using social structure. In *Proceedings of the Joint Statistical Meetings*. Washington, DC.
- Federal Bureau of Investigation (1999). Crime in the United States.
- Office of Advocacy, U.S. Small Business Administration (1997). Characteristics of small business employees and owners.
- SALGANIK, M. J., MELLO, M. B., ADBO, A. H., BERTONI, N., FATZIO, D. and BASTOS, F. I. (2011). The game of contacts: Estimating the social visibility of groups. *Social Networks* **33** 70–78.
- SHELLY, G., BERNARD, H., KILLWORTH, P., JOHNSEN, E. and MCCARTY, C. (1995). Who knows you HIV status? What HIV+ patients and their network members know about each other. *Social Networks* **17** 189–217.
- UNAIDS (2003). Estimating the size of popualtions at risk for HIV 03.36E. UNAIDS, Geneva.
- ZHENG, T., SALGANIK, M. J. and GELMAN, A. (2006). How many people do you know in prison?: Using overdispersion in count data to estimate social structure in networks. *J. Amer. Statist. Assoc.* **101** 409–423. [MR2256163](#)

DEPARTMENT OF STATISTICS
 DEPARTMENT OF SOCIOLOGY
 CENTER FOR STATISTICS AND THE SOCIAL SCIENCES
 UNIVERSITY OF WASHINGTON
 Box 354320 SEATTLE, WASHINGTON 98195
 USA
 E-MAIL: tylermc@u.washington.edu

DEPARTMENT OF STATISTICS
 COLUMBIA UNIVERSITY
 1255 AMSTERDAM AVE.
 NEW YORK, NEW YORK 10027
 USA
 E-MAIL: tzheng@stat.columbia.edu
 URL: <http://www.stat.columbia.edu/~tzheng>